

# TPRM: Tensor partition regression models with applications in imaging biomarker detection

Michelle F. Miranda\*

The Research, Innovation and Dissemination Center for  
Neuromathematics, Universidade de São Paulo

and

Hongtu Zhu <sup>†</sup>

Joseph G. Ibrahim <sup>‡</sup>

Department of Biostatistics, University of North Carolina at Chapel Hill

and

for the Alzheimer's Disease Neuroimaging Initiative<sup>§</sup>

May 21, 2015

## Abstract

Many neuroimaging studies have collected ultra-high dimensional imaging data in order to identify imaging biomarkers that are related to normal biological processes, diseases, and the response to treatment, among many others. These imaging data are

---

\*Dr. Miranda's research was supported by grant #2013/ 07699-0 , S. Paulo Research Foundation.

<sup>†</sup>Dr. Zhu was supported by NIH grants 1UL1TR001111 and MH086633, and NSF Grants SES-1357666 and DMS-1407655.

<sup>‡</sup>Dr. Ibrahim's research was partially supported by NIH grants #GM 70335 and P01CA142538.

<sup>§</sup>Data used in preparation of this article were obtained from the Alzheimer's Disease Neuroimaging Initiative (ADNI) database (adni.loni.usc.edu). As such, the investigators within the ADNI contributed to the design and implementation of ADNI and/or provided data but did not participate in analysis or writing of this report. A complete listing of ADNI investigators can be found at: [http://adni.loni.usc.edu/wp-content/uploads/how\\_to\\_apply/ADNI\\_Acknowledgement\\_List.pdf](http://adni.loni.usc.edu/wp-content/uploads/how_to_apply/ADNI_Acknowledgement_List.pdf).

often represented in the form of a multi-dimensional array, called a tensor. Existing statistical methods are insufficient for analysis of these tensor data due to their ultra-high dimensionality as well as their complex structure. The aim of this paper is develop a tensor partition regression modeling (TPRM) framework to establish an association between low-dimensional clinical outcomes and ultra-high dimensional tensor covariates. Our TPRM is a hierarchical model and efficiently integrates four components: (i) a partition model; (ii) a canonical polyadic decomposition model; (iii) a factor model; and (iv) a generalized linear model. Under this framework, ultra-high dimensionality is not only reduced to a manageable level, resulting in efficient estimation, but also prediction accuracy is optimized to search for informative sub-tensors. Posterior computation proceeds via an efficient Markov chain Monte Carlo algorithm. Simulation shows that TPRM outperforms several other competing methods. We apply TPRM to predict disease status (Alzheimer versus control) by using structural magnetic resonance imaging data obtained from Alzheimer’s Disease Neuroimaging Initiative (ADNI) study.

*Keywords:* Big data; ADNI; Bayesian hierarchical model; Tensor decomposition; Tensor regression; Chib augmentation method; MCMC.

# 1 Introduction

Many neuroimaging studies have collected ultra-high dimensional imaging data in order to identify imaging biomarkers that are related to normal biological processes, diseases, and the response to treatment, among many others. The imaging data provided by these studies are often represented in the form of a multi-dimensional array, called a tensor. Existing statistical methods are insufficient for analysis of these tensor data due to their ultra-high dimensionality as well as their complex structure.

The aim of this paper is to develop a novel tensor partition regression modeling (TPRM) framework to use high-dimensional imaging data, denoted by  $\mathbf{x}$ , to predict a scalar response, denoted by  $y$ . The scalar response  $y$  may include cognitive outcome, disease status, and the early onset of disease, among others. In various neuroimaging studies, imaging data are often measured at a large number of grid points in a three (or higher) dimensional space and have a multi-dimensional tensor structure. Without loss of generality, we use  $\mathbf{x} = (x_{j_1 \dots j_D}) \in R^{J_1 \times \dots \times J_D}$  to denote an order  $D$  tensor, where  $D \geq 2$ . Vectorizing  $\mathbf{x}$  leads to a  $(\prod_{k=1}^D J_k) \times 1$  vector. Examples of  $\mathbf{x}$  include magnetic resonance imaging (MRI), diffusion tensor imaging (DTI), and positron emission tomography (PET), among many others. These advanced medical imaging technologies are essential to understanding the neural development of neuropsychiatric and neurodegenerative disorders and normal brain development.

Although a large family of regression methods have been developed for supervised learning (Hastie et al., 2009; Breiman et al., 1984; Friedman, 1991; Zhang and Singer, 2010), their computability and theoretical guarantee are compromised by this ultra-high dimensionality of imaging data. The first set of promising solutions is high-dimensional sparse regression (HSR) models, which often take high-dimensional imaging data as unstructured predictors. A key assumption of HSR is its sparse solutions. HSRs not only suffer from diverging spectra and noise accumulation in ultra-high dimensional feature spaces (Fan and Fan, 2008; Bickel and Levina, 2004), but also their sparse solutions may lack biological interpretation in neuroimaging studies. Moreover, standard HSRs ignore the inherent spatial structure of the image that possesses a wealth of spatial information, such as spatial correlation and spatial smoothness. To address some limitations of HSRs, a family of tensor regression models has been developed to preserve the spatial structure of imaging tensor data, while achieving substantial dimensional reduction (Zhou et al., 2013).

The second set of solutions adopts functional linear regression (FLR) approaches, which treat imaging data as functional predictors. However, since most existing FLR models focus on one dimensional curves (Müller and Yao, 2008; Ramsay and Silverman, 2005), general-

izations to two and higher dimensional images is far from trivial and requires substantial research (Reiss and Ogden, 2010). Most estimation approaches of FLR approximate the coefficient function of such functional regression models as a linear combination of a set of fixed (or data-driven) basis functions. For instance, most estimation methods of FLR based on the fixed basis functions (e.g., tensor product wavelet) are required to solve an ultra-high dimensional optimization problem and suffer the same limitations as those of HSR.

The third set of solutions usually integrates supervised (or unsupervised) dimension reduction techniques with various standard regression models. Given the ultra-high dimension of imaging data, however, it is imperative to use some dimension reduction methods to extract and select ‘low-dimensional’ important features, while eliminating most redundant features (Johnstone and Lu, 2009; Bair et al., 2006; Fan and Fan, 2008; Tibshirani et al., 2002; Krishnan et al., 2011). Most of these methods first carry out an unsupervised dimension reduction step, often by principal component analysis (PCA), and then fit a regression model based on the top principal components (Caffo et al., 2010). Recently, for ultra-high tensor data, unsupervised higher order tensor decompositions (e.g. parallel factor analysis and Tucker) have been extensively proposed to extract important information of neuroimaging data (Martinez et al., 2004; Beckmann and Smith, 2005; Zhou et al., 2013). Although it is intuitive and easy to implement such methods, it is well known that the features extracted from PCA and Tucker can be irrelevant to the response.

In this paper, we develop a novel TPRM to establish an association between imaging tensor predictors and clinical outcomes. Our TPRM is a hierarchical model with four components: (i) a partition model that divides the high-dimensional tensor covariates into sub-tensor covariates; (ii) a canonical polyadic decomposition model that reduces the sub-tensor covariates to low-dimensional feature vectors; (iii) a generalized linear model that uses the feature vectors to predict clinical outcomes; (iv) a sparse inducing normal mixture prior is used to select informative feature vectors. Although the four components of TPRM have been independently developed/used in different settings, the key novelty of TPRM lies in the integration of (i)-(iv) into a single framework for imaging prediction. In particular, the first two components (i) and (ii) are designed to specifically address the three key features of neuroimaging data: relatively low signal to noise ratio, spatially clustered effect regions, and the tensor structure of imaging data. The neuroimaging data are often very noisy, while the ‘activated’ (or ‘effect’) brain regions associated with the response are usually clustered together and their size can be very small. In contrast, a crucial assumption for the success of most matrix/array decomposition methods (e.g., singular value decomposition) is that the leading components obtained from these decomposition

methods capture the most important feature of a multi-dimensional array. Under TPRM, the ultra-high dimensionality of imaging data is dramatically reduced by using the partition model. For instance, let's consider a standard  $256 \times 256 \times 256$  3D array with 16,777,216 voxels, and its partition model with  $32^3 = 32,768$  sub-arrays with size  $8 \times 8 \times 8$ . If we reduce each  $8 \times 8 \times 8$  into a small number of components by using component (ii), then the total number of reduced features is around  $O(10^4)$ . We can further increase the size of each subarray in order to reduce the size of neuroimaging data to a manageable level, resulting in efficient estimation.

The rest of the article is organized as it follows. In Section 2, we introduce TPRM, the priors, and a Bayesian estimation procedure. In Section 3, we use simulated data to compare the Bayesian decomposition with several competing methods. In Section 4, we apply our model to the ADNI data set. In Section 5, we present some concluding remarks.

## 2 Methodology

### 2.1 Preliminaries

We review several basic facts of tensors (Kolda and Bader, 2009). A **tensor**  $\mathbf{x} = (x_{j_1 \dots j_D})$  is a multidimensional array, whose order  $D$  is determined by its dimension. For instance, a vector is a tensor of order 1 and a matrix is a tensor of order 2.

The **inner product** between two tensors  $\mathcal{X} = (x_{j_1 \dots j_D})$  and  $\mathcal{Y} = (y_{j_1 \dots j_D})$  in  $\Re^{J_1 \times \dots \times J_D}$  is the sum of the product of their entries given by

$$\langle \mathcal{X}, \mathcal{Y} \rangle = \sum_{j_1=1}^{J_1} \dots \sum_{j_D=1}^{J_D} x_{j_1 \dots j_D} y_{j_1 \dots j_D}.$$

The **outer product** between two vectors  $\mathbf{a}^{(1)} = (a_{j_1}^{(1)}) \in \Re^{J_1}$  and  $\mathbf{a}^{(2)} = (a_{j_2}^{(2)}) \in \Re^{J_2}$  is a matrix  $M = (m_{j_1 j_2})$  of size  $J_1 \times J_2$  with entries  $m_{j_1 j_2} = a_{j_1}^{(1)} a_{j_2}^{(2)}$ . A tensor  $\mathcal{X} \in \Re^{J_1 \times \dots \times J_D}$  is a *rank one tensor* if it can be written as an outer product of  $D$  vectors such that  $\mathcal{X} = \mathbf{a}^{(1)} \circ \mathbf{a}^{(2)} \dots \circ \mathbf{a}^{(D)}$ , where  $\mathbf{a}^{(k)} \in \Re^{J_k}$  for  $k = 1, \dots, D$ . Moreover, the parallel factor analysis, also known as PARAFAC or **CP decomposition**, factorizes a tensor into a sum of rank-one tensors such that

$$\mathcal{X} \approx \sum_{r=1}^R \lambda_r \mathbf{a}_r^{(1)} \circ \mathbf{a}_r^{(2)} \circ \dots \circ \mathbf{a}_r^{(D)}, \quad (1)$$

where  $\mathbf{a}_r^{(k)} = (a_{j_k r}^{(k)}) \in \Re^{J_k}$  for  $k = 1, \dots, D$  and  $r = 1, \dots, R$ . See Figure 1 for an illustration

of a 3D array.




Figure 1: Figure copied from (Kolda and Bader, 2009). Panel (a) illustrates the CP decomposition of a three way array as a sum of  $R$  components of rank-one tensors, i.e.  $\mathcal{X} \approx \sum_{r=1}^R \mathbf{a}_r \circ \mathbf{b}_r \circ \mathbf{c}_r$ .

We need the following notation throughout the paper. Suppose that we observe data  $\{(y_i, \mathcal{X}_i, \mathbf{z}_i) : i = 1, \dots, N\}$  from  $n$  subjects, where the  $\mathcal{X}_i$ 's are tensor imaging data,  $\mathbf{z}_i$  is a  $p_z \times 1$  vector of scalar covariates, and  $y_i$  is a scalar response, such as diagnostic status or clinical outcome. If we concatenate all  $D$  dimensional tensor  $\mathcal{X}_i$ s into a  $(D+1)$  dimensional tensor  $\tilde{\mathcal{X}} = \{\mathcal{X}_i, i = 1, \dots, N\} = (x_{j_1, \dots, j_D, i})$ , we consider the CP decomposition of  $\tilde{\mathcal{X}}$  as follows:

$$\tilde{\mathcal{X}} = \|\Lambda; \mathbf{A}^{(1)}, \dots, \mathbf{A}^{(D)}, \mathbf{L}\| \quad \text{or} \quad x_{j_1, \dots, j_D, i} = \sum_{r=1}^R \lambda_r a_{j_1 r}^{(1)} a_{j_2 r}^{(2)} \dots a_{j_D r}^{(D)} l_{ir}. \quad (2)$$

where  $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_R)$ ,  $\mathbf{A}^{(d)} = [\mathbf{a}_1^{(d)} \mathbf{a}_2^{(d)} \dots \mathbf{a}_R^{(d)}]$  for  $d = 1, \dots, D$ , and  $\mathbf{L} = (l_{ir})$ . The matrices  $\mathbf{A}^{(d)}$ 's and  $\mathbf{L}$  are called factor matrices.

## 2.2 Tensor Partition Regression Models

Our interest is to develop TPRM for establishing the association between responses  $y$  and their corresponding imaging covariates  $\mathcal{X}$  and clinical covariates  $\mathbf{Z}$ . The first component of TPRM is a partition model that divides the high-dimensional tensor  $\tilde{\mathcal{X}}$  into  $S$  disjoint sub-tensor covariates  $\tilde{\mathcal{X}}^{(s)}$ , that is

$$\tilde{\mathcal{X}} = \cup_{s=1}^S \tilde{\mathcal{X}}^{(s)} \quad \text{and} \quad \tilde{\mathcal{X}}^{(s)} \cap \tilde{\mathcal{X}}^{(s')} = \emptyset. \quad (3)$$

Although the size of  $\tilde{\mathcal{X}}^{(s)}$  can vary across  $s$ , it is assumed that without loss of generality,  $\tilde{\mathcal{X}}^{(s)} \in \mathcal{R}^{p_1 \times \dots \times p_D}$  and the size of  $\tilde{\mathcal{X}}^{(s)}$  is homogeneous such that  $S = \prod_{k=1}^D (J_k/p_k)$ .

The second component of TPRM is a canonical polyadic decomposition model that

reduces the sub-tensor covariates  $\tilde{\mathcal{X}}^{(s)}$  to low-dimensional feature vectors. Specifically, it is assumed that for each  $s$ , we have

$$\tilde{\mathcal{X}}^{(s)} = \|\mathbf{A}_s; \mathbf{A}_s^{(1)}, \mathbf{A}_s^{(2)}, \dots, \mathbf{A}_s^{(D)}, \mathbf{L}_s\| + \mathcal{E}^{(s)}, \quad (4)$$

where  $\mathbf{A}_s = \text{diag}(\lambda_1^{(s)}, \dots, \lambda_R^{(s)})$  consists of the weights for each rank of the decomposition in (4),  $\mathbf{A}_s^{(d)} \in \mathbb{R}^{p_d \times R}$  are the factor matrices along the  $d$ -th dimension of  $\mathcal{X}$ , and  $\mathbf{L}_s \in \mathbb{R}^{N \times R}$  is the factor matrix along the subject dimension. It is assumed that the elements of  $\mathcal{E}^{(s)} = (e_{j_1 \dots j_{D+1}}^{(s)})$  are measurement errors and  $e_{j_1 \dots j_{D+1}}^{(s)} \sim N(0, (\tau^{(s)})^{-1})$ . The elements of  $\mathbf{L}_s$  capture the major variation in  $\mathcal{X}^{(s)}$  due to subject differences, while the common structure among the subjects is absorbed into the factor matrices  $\mathbf{A}_s^{(d)}$  for  $d = 1, \dots, D$  (Kolda and Bader, 2009).

There are two key advantages of using (3) and (4). First, the use of the partition model (3) allows us to concentrate on the most important local features of each sub-tensor, instead of the major variation of the whole image, which may be unassociated with the response of interest. In many applications, although the effect regions associated with responses may be relatively small compared with the whole image, their size can be comparable with that of each sub-tensor. Therefore, one can extract more informative features associated with the response with a high probability. Second, the use of the canonical polyadic decomposition model (4) can substantially reduce the dimension of the original imaging data. Recall the discussions in Section 1 that the use of  $8 \times 8 \times 8$  sub-tensors can substantially reduce imaging size at a scale of  $O(10^3)$ .

The third component of TPRM is a factor model that decomposes  $\mathbf{L} = [\mathbf{L}_1, \dots, \mathbf{L}_S]$  into the product

$$\mathbf{L} = \mathbf{G}\mathbf{D} + \mathbf{\Psi}, \quad (5)$$

where each row of the matrix  $\mathbf{G}$  is a  $K \times 1$  vector of common unobserved (latent) factors  $\mathbf{g}_i$ ;  $\mathbf{D} \in \mathbb{R}^{K \times P_L}$  correspond to the matrix of  $K$  latent basis functions used to represent  $\mathbf{L}$ ;  $\mathbf{\Psi}$  is a matrix representing idiosyncratic errors. The proposed factor model not only reduces the feature vector into a manageable level but it also deals with the multicollinearity that may exist between features of the same partition.

The fourth component of TPRM is a generalized linear model that links scalar responses  $y_i$  and their corresponding reduced imaging features  $\mathbf{g}_i$  and clinical covariates  $\mathbf{z}_i$ . Specifically,  $y_i$  given  $\mathbf{g}_i$  and  $\mathbf{z}_i$  follows an exponential family distribution with density given by

$$f(y_i|\boldsymbol{\theta}_i) = m(y_i) \exp\{\eta(\boldsymbol{\theta}_i)T(y_i) - a(\boldsymbol{\theta}_i)\}, \quad (6)$$

where  $m(\cdot)$ ,  $\eta(\cdot)$ ,  $T(\cdot)$ , and  $a(\cdot)$  are pre-specified functions. Moreover, it is assumed that  $\mu_i = E(y_i|\mathbf{g}_i, \mathbf{z}_i)$  satisfies

$$h(\mu_i) = \mathbf{z}_i^T \boldsymbol{\gamma} + \mathbf{g}_i^T \mathbf{b}, \quad (7)$$

where  $\boldsymbol{\gamma}$  and  $\mathbf{b}$  are coefficient vectors associated with  $\mathbf{z}_i$  and  $\mathbf{g}_i$ , respectively and  $h(\cdot)$  is a link function.

## 2.3 Prior Distributions

We consider the priors on the elements of  $b_k$ . Bimodal sparsity promoting priors are key elements to perform variable selection and have been the subject of extensive research (Mayrink and Lucas, 2013; George and McCulloch, 1993, 1997). We assume the following hierarchy:

$$\begin{aligned} b_k | \delta_k, \sigma^2 &\sim (1 - \delta_k)F(b_k) + \delta_k N(0, \sigma^2), \\ \delta_k | \pi &\sim \text{Bernoulli}(\pi) \quad \text{and} \quad \pi \sim \text{Beta}(\alpha_{0\pi}, \alpha_{1\pi}), \end{aligned} \quad (8)$$

where  $F(\cdot)$  is a pre-specified probability distribution. A common choice of  $F(\cdot)$  is a degenerate distribution at 0, leading to what is called the ‘spike and slab’ prior (Mitchell and Beauchamp, 1988). A different approach is to consider  $F = N(0, \epsilon)$  with a very small  $\epsilon$  instead of putting a probability mass on  $b_k = 0$ . Thus, the  $b_k$ ’s are assumed to come from a mixture of two normal distributions. In this case, the hyperparameter  $\sigma^2$  should be large enough to give support to values of the coefficients that are substantively different from 0, but not so large that unrealistic values of  $b_k$  are supported. In this article, we opt for the latter approach.

The probability  $\pi$  determines whether a particular component of  $\mathbf{g}_i$  is informative for predicting  $y$ . A common choice for its prior is a non-informative distribution with  $\alpha_{0\pi} = \alpha_{1\pi} = 1$ . However, this choice of the hyperparameters implies that its posterior mean is restricted to the interval  $[1/3, 2/3]$ , a undesirable feature in variable selection. To fix this, we choose a ‘bathtub’ shaped beta distribution, since a prior concentrating most of its mass in the extremes of the interval  $(0, 1)$  is evidently more suitable for variable selection (Gonçalves et al., 2013).

We consider the priors on the elements of  $A_{(s)r}^{(d)}$ ,  $\mathbf{l}_r^{(s)}$ ,  $\tau^{(s)}$ , and  $\lambda_r^{(s)}$ . For  $d = 1, \dots, D$  and  $r = 1, \dots, R$ , we assume

$$\begin{aligned} A_{(s)r}^{(d)} &\sim N(0, p_d^{-1} \mathbf{I}_{p_d}), \quad \mathbf{l}_r^{(s)} \sim N(0, \mathbf{I}_N), \\ \tau^{(s)} &\sim \text{Gamma}(\nu_{0\tau}, \nu_{1\tau}), \quad \text{and} \quad \lambda_r^{(s)} \sim N(0, \kappa^{-1}). \end{aligned} \quad (9)$$



For  $(\gamma, v)$ , it is assumed that  $\gamma \sim N(\gamma^*, v^{-1} \mathbf{I}_q)$  and  $v \sim \text{Gamma}(\nu_{0v}, \nu_{1v})$ .

For the elements of the factor model, we let  $\mathbf{G} = \{\mathbf{g}_1, \dots, \mathbf{g}_K\}$ ,  $\mathbf{D} = \{\mathbf{d}_1, \dots, \mathbf{d}_{P_L}\}$  and the elements of  $\Psi$ ,  $\psi_{ij} \sim N(0, \tau_\psi^{-1})$  for  $i = 1, \dots, n$  and  $j = 1, \dots, P_L$ . For each fixed  $k = 1, \dots, K$ , it is assumed that

$$\mathbf{g}_k \sim N(0, n^{-1} \mathbf{I}_n), \quad \mathbf{d}_k \sim N(0, K^{-1} \mathbf{I}_K), \quad \text{and} \quad \tau_\psi \sim \text{Gamma}(\beta_{0\tau_\psi}, \beta_{1\tau_\psi}),$$

where  $\mathbf{I}_k$  be a  $k \times k$  identity matrix. When  $p_d$  is large, the columns of the factor matrix  $\mathbf{A}_{(s)r}^{(d)}$  are approximately orthogonal, which is consistent with their role in the decomposition (1). When  $n$  and  $K$  are large, the matrices  $\mathbf{G}$  and  $\mathbf{D}$  are approximately orthogonal to each other. However, we only impose that the columns of these matrices span the space of the principal vectors, without explicitly requiring orthonormality, which leads substantially computational efficiency (Xinghao Ding and Carin, 2011).

## 2.4 Posterior Inference

Let  $\boldsymbol{\theta} = \{\mathbf{A}^{(1)} \dots, \mathbf{A}^{(D)}, \mathbf{L}, \mathbf{\Lambda}, \boldsymbol{\tau}, \mathbf{G}, \mathbf{D}, \tau_\psi, \mathbf{B}, \boldsymbol{\delta}, \pi, \gamma, v\}$ , where  $\mathbf{A}^{(d)} = [\mathbf{A}_{(1)}^{(d)}, \dots, \mathbf{A}_{(S)}^{(d)}]$ ,  $\mathbf{L} = [\mathbf{L}_1, \dots, \mathbf{L}_S]$ ,  $\mathbf{\Lambda} = [\mathbf{\Lambda}_1, \dots, \mathbf{\Lambda}_S]$ , and  $\boldsymbol{\tau} = [\tau^{(1)}, \dots, \tau^{(S)}]$ . A Gibbs sampler algorithm is used to generate a sequence of random observations from the joint posterior distribution given by

$$\begin{aligned} p(\boldsymbol{\theta} | \mathcal{X}, \mathbf{y}) &\propto p(\mathbf{y} | \boldsymbol{\theta}, \mathcal{X}) p(\mathbf{A}^{(1)} \dots, \mathbf{A}^{(D)}, \mathbf{L}, \mathbf{\Lambda}, \boldsymbol{\tau} | \mathcal{X}) \\ &\quad p(\mathbf{G}, \mathbf{D}, \tau_\psi | \mathbf{L}) p(\mathbf{B} | \mathbf{y}, \mathbf{G}, \boldsymbol{\delta}) p(\boldsymbol{\delta} | \pi) p(\pi) p(\gamma | v) p(v). \end{aligned} \quad (10)$$

The Gibbs sampler essentially involves sampling from a series of conditional distributions, while each of the modeling components is updated in turn. The detailed sampling algorithm is described on Appendix B.

## 3 Simulation Study

We carried out three sets of simulations to examine the finite-sample performance of TPRM and its associated Gibbs sampler algorithm.

### 3.1 Bayesian tensor decomposition

The goals of the first set of simulations are (i) to compare the proposed Bayesian tensor decomposition method with the alternating least squares method, (ii) to investigate

how different choices of the rank  $R$  impact the tensor decomposition for distinct image modalities; and (iii) to assess the importance of the partition model. We considered 3 different imaging data sets (or tensors) including (I-1) a diffusion tensor image (DTI) of size  $90 \times 96 \times 96$ , (I-2) a white matter RAVENS map image of size  $99 \times 99 \times 70$ , and (I-3) a T1-weighted MRI image of size  $64 \times 108 \times 99$ . We fitted models (3) and (4) to the three image tensors and decomposed each of them with  $R = 5, 10$ , and  $20$ . We consider 27 partitions of size  $30 \times 30 \times 32$  for the DTI image, 18 partitions of size  $33 \times 33 \times 35$  for the RAVENS map, and 24 partitions of size  $32 \times 27 \times 33$  for the T1 image, respectively. The hyperparameters were chosen to reflect non-informative priors and are set as  $\nu_{0\tau} = 1$ ,  $\nu_{1\tau} = 10^{-2}$ , and  $\kappa = 10^{-6}$ .

We run steps (a.1) – (a.4) of the Gibbs sampler algorithm in Section 2.4 for 5,000 iterations. The efficiency of the proposed algorithm is observed through trace plots for 9 random voxels. Figure 2 shows the results for the reconstructed white matter RAVENS map decomposed with  $R = 20$ . The proposed algorithm is efficient and has fast convergence. At each iteration, we computed the quantity  $\mathcal{I} = \sum_{s=1}^S \|\Lambda_s; \mathbf{A}_s^{(1)}, \mathbf{A}_s^{(2)}, \mathbf{A}_s^{(3)}\|$  for each rank and each partition. Subsequently, we computed the reconstructed image, defined as  $\hat{\mathcal{X}}$ , and the posterior mean estimate of  $\mathcal{I}$  after a burn-in sample of 3,000. For each reconstructed image  $\hat{\mathcal{X}}$ , we computed its root mean squared error,  $\text{RMSE} = \|\hat{\mathcal{X}} - \mathcal{X}\|_2 / \sqrt{J_1 J_2 J_3}$ .

We consider the non-partition model and compare the Bayesian method and the standard alternating least squares method (Kolda and Bader, 2009). The results are shown in Table 1. Figure 3 shows an axial slice of the original images and the reconstructed images for ranks  $R = 5, 10$ , and  $20$  as  $S = 1$ . For the images considered in this study, the Bayesian decomposition gives a smaller RMSE for all cases. As expected, the higher the rank, the smaller the reconstruction error.

## 3.2 A 2-dimensional image example

The goals of the second set of simulations are to assess whether TPRM is able to capture regions of interest, that significantly differ between two groups, in a 2-dimensional phantom and to compare TRPM with the functional principal components model (fPCA). We generate a data set  $\{(y_i, \mathcal{X}_i) : i = 1, \dots, n\}$  with  $n = 200$  according to

$$y_i \sim \text{Bernoulli}(0.5) \quad \text{and} \quad \mathcal{X}_i = X_0(y_i) + \mathcal{E}_i,$$

where  $\mathcal{X}_i$ ,  $\mathcal{E}_i = (\epsilon_{ij_1 j_2})$ , and  $X_0(y_i)$  are  $32 \times 32$  matrices and  $X_0(1)$  and  $X_0(1) - X_0(0)$  are, respectively, shown in panels (a) and (b) of Figure 4. We independently generated  $\epsilon_{ij_1 j_2}$




Figure 2: Trace plots in 9 randomly chosen voxels in the white matter RAVENS map by using Bayesian tensor decomposition with  $R = 20$ . The trace plots indicate that the Markov chain converges after around 1000 iterations.

Table 1: Root mean squared error for 3 different image modalities. The Bayesian decomposition outperforms the alternating least squares in each scenario. There is a smaller error measurement with an increase of the rank  $R$ .

		T1-weighted	WM RAVENS	DTI
R=5	BayesianCP	45.3191	1.5853	3.1656e-004
	ALS	45.3636	1.6013	3.2506e-004
	Partition	37.3712	1.2178	2.0929e-004
R=10	BayesianCP	41.7018	1.4382	2.7367e-004
	ALS	42.4350	1.4533	2.8247e-004
	Partition	31.3836	1.0186	1.5748e-004
R=20	BayesianCP	37.1796	1.2885	2.2911e-004
	ALS	38.3166	1.3166	2.3676e-004
	Partition	25.1574	0.8085	1.1349e-004

from a  $N(0, 5^2)$  generator for  $j_1, j_2 = 1, \dots, 32$ . Panel (c) in Figure 4 shows a generated 2D image from a random subject in group 1, which is almost indistinguishable from random noise. The hyperparameters in Section 2.3 are chosen to reflect non-informative priors with  $\nu_{0\tau} = 1$ ,  $\nu_{1\tau} = 10^{-4}$ ,  $\sigma^2 = 10^4$ , and  $\kappa = 10^{-4}$ .

We applied two TPRMs with  $S = 1$  (no-partition model) and  $S = 16$  to the simulated data set. We compared the two TPRMs with the functional principal components model (fPCA), in which we learned the basis functions in the first stage and then included the




Figure 3: Simulation 1 for Bayesian tensor decomposition results. In each panel, the image on the left represents an axial slice of the original image, with top panel representing the non-partition model and bottom panels the partition model. In addition, from left to right we have the decomposed images for ranks  $\mathcal{R} = 5, 10, 20$ , respectively. Panel (a): DTI image; panel (b): white matter RAVENS map; and panel (c): T1-weighted image.

top  $R$  most important principal components as covariates in a probit regression in the second stage. We set  $R = 8$  for all three models, and skip the factor model in (5) since the number of features is manageable in this simulation. For TPRMs, we run the Gibbs sampler algorithm for 5000 iterations with a burn-in period of 3000 iterations.

We also computed the Bayesian estimate of  $\mathcal{P} = \|\mathbf{\Lambda}; \mathbf{A}^{(1)}, \mathbf{A}^{(2)}, \mathbf{B}\|$  by using MCMC samples. The estimated quantity  $\mathcal{P}$  represents a projection of the group differences into the image space. Furthermore, we used MCMC samples to construct credible intervals for  $\mathcal{P}$  in the imaging space. This quantity is extremely important in neuroimaging studies since it allows us to precisely identify significant locations in the brain that are associated with the response variable. Panels (c), (d) and (e) of Figure 4 are the posterior mean estimates of  $\mathcal{P}$  for the fPCA model, TPRM with  $S = 1$ , and TPRM with  $S = 16$ , respectively. Panel (f) of Figure 4 shows the 95% credible interval for TPRMs with  $S = 16$ . The result reveals that the proposed model closely recovers the true underlying location where differences between both groups exist.

### 3.3 A 3D image example

The goal of this set of simulations is to examine the classification performance of the partition model in the 3D imaging setting. To mimic real data, we consider scenarios where we have extremely noisy images. Our goal is to compare 3 models whose main difference is the way the features are extracted: (i) functional principal component model (fpca); (ii) tensor alternating least squares (tals); and (iii) partition model with tensor decomposition and principal components on the extracted features (pmted). We simulated the three-dimensional image covariates  $\mathcal{X}_i(y_i)$  as follows:

$$\mathcal{X}_i(y_i) = \mathcal{G}_0 + c_0 y_i \mathcal{X}_0 + \mathcal{E}_i,$$

for  $i = 1, \dots, 200$ ,  $\mathcal{G}_0 \in \mathbb{R}^{64 \times 64 \times 50}$  is a fixed brain template with values ranging from 0 to 250,  $c_0 = 50, 65$ , and the elements of the tensor  $\mathcal{E}_i \in \mathbb{R}^{64 \times 64 \times 50}$  were independently generated from a  $N(0, 70^2)$  generator. Moreover, we set  $\mathcal{X}_0 = \|\mathbf{1}; \mathbf{A}^{(1)}, \mathbf{A}^{(2)}, \mathbf{A}^{(3)}\|$ , where  $\mathbf{A}_0^{(1)} \in \mathbb{R}^{64 \times 4}$ ,  $\mathbf{A}_0^{(2)} \in \mathbb{R}^{64 \times 4}$ , and  $\mathbf{A}_0^{(3)} \in \mathbb{R}^{50 \times 4}$ , and  $\mathbf{A}_0^{(d)}$  are matrices whose  $(c_d + j)$ -th element of each column is equal to  $\sin(j\pi/14)$ , with  $c_d$  indicating the position at the  $d$ -th coordinate.

For each value of  $c_0$ , we consider  $S = 32$  and run the probit model for 100 generated datasets. For each run we compute the model accuracy defined as  $1 - (\sum_i y_i - I(\Phi(\mathbf{X}_i \hat{\mathbf{b}}) > 0.5))/N$ , where  $I(A)$ , is the indicator function of the set  $A$ . The results are summarized

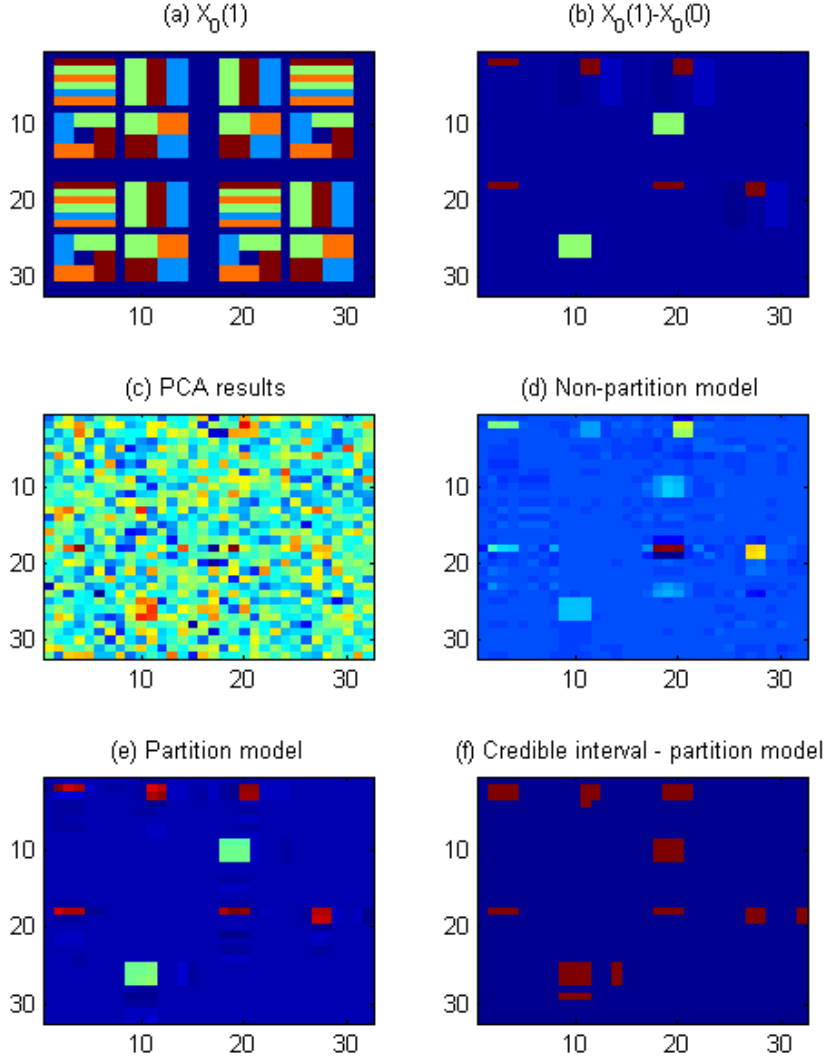


Figure 4: Results of the 2-D imaging example: (a):  $\mathcal{X}_0(1)$ ; (b):  $\mathcal{X}_0(1) - \mathcal{X}_0(0)$ ; (c): the estimated projection  $\mathcal{P}$  for the fPCA model. Panels (d) and (e) are the posterior mean of the quantity  $\mathcal{P} = \|\Lambda; \mathbf{A}^{(1)}, \mathbf{A}^{(2)}, \mathbf{B}\|$  for the non-partition model BTRM with  $S = 1$  and for BTRM with  $S = 16$ , respectively. Panel (f) is the 95% credible interval of  $\mathcal{P}$  for BTRM( $S = 16$ ) revealing the true underlying location, where differences between both groups exist.

in Table 2. Further, we consider one generated dataset for each scenario and run a 10-fold cross-validation procedure. Each training set has 180 subjects and 20 testing subjects. The prediction accuracy is then computed on the test set and a summary with the results is also shown in Table 2. Our partition model outperforms the tensor model and the fpca for both scenarios.

Table 2: Mean model accuracy and mean prediction accuracy. The partition model outperforms the tensor model and the fpca for both measurements and both scenarios. As expected, the model accuracy is higher than the prediction accuracy.

		fpca	tals	pmtl
Monte Carlo	$cc = 50$	0.7750	0.7100	<b>0.8650</b>
	$cc = 65$	0.8650	0.745	<b>0.9318</b>
10-fold CV	$cc = 50$	0.5900	0.5050	<b>0.7500</b>
	$cc = 65$	0.6400	0.5400	<b>0.7800</b>

## 4 Real data analysis

“Data used in the preparation of this article were obtained from the Alzheimer’s Disease Neuroimaging Initiative (ADNI) database ([adni.loni.usc.edu](http://adni.loni.usc.edu)). The ADNI was launched in 2003 by the National Institute on Aging (NIA), the National Institute of Biomedical Imaging and Bioengineering (NIBIB), the Food and Drug Administration (FDA), private pharmaceutical companies and non-profit organizations, as a \$60 million, 5-year public private partnership. The primary goal of ADNI has been to test whether serial magnetic resonance imaging (MRI), positron emission tomography (PET), other biological markers, and clinical and neuropsychological assessment can be combined to measure the progression of mild cognitive impairment (MCI) and early Alzheimer’s disease (AD). Determination of sensitive and specific markers of very early AD progression is intended to aid researchers and clinicians to develop new treatments and monitor their effectiveness, as well as lessen the time and cost of clinical trials. The Principal Investigator of this initiative is Michael W. Weiner, MD, VA Medical Center and University of California, San Francisco. ADNI is the result of efforts of many coinvestigators from a broad range of academic institutions and private corporations, and subjects have been recruited from over 50 sites across the U.S. and Canada. The initial goal of ADNI was to recruit 800 subjects but ADNI has been followed by ADNI-GO and ADNI-2. To date these three protocols have recruited over 1500 adults, ages 55 to 90, to participate in the research, consisting of cognitively normal

older individuals, people with early or late MCI, and people with early AD. The follow up duration of each group is specified in the protocols for ADNI-1, ADNI-2 and ADNI-GO. Subjects originally recruited for ADNI-1 and ADNI-GO had the option to be followed in ADNI-2. For up-to-date information, see [www.adni-info.org](http://www.adni-info.org).”

We applied the proposed model to the anatomical MRI data collected at the baseline of ADNI. We considered 402 MRI scans from ADNI1, 181 of them were diagnosed with AD, and 221 healthy controls. These scans were performed on a 1.5 T MRI scanners using a sagittal MPRAGE sequence and the typical protocol includes the following parameters: repetition time (TR) = 2400 ms, inversion time (TI) = 1000 ms, flip angle = 80, and field of view (FOV) = 24 cm with a  $256 \times 256 \times 170$  mm<sup>3</sup> acquisition matrix in the x, y, and z dimensions, which yields a voxel size of  $1.25 \times 1.26 \times 1.2$  mm<sup>3</sup> (Huang et al., 2014).

The T1-weighted images were processed using HAMMER (Hierarchical Attribute Matching Mechanism for Elastic Registration), a free pipeline. The processing steps include skull and cerebellum removal, followed by tissue segmentation to identify the regions of white matter (WM), gray matter (GM) and cerebrospinal fluid (CSF). Then, registration was performed to warp the subject to the space of the Jacob template (size  $256 \times 256 \times 256$  mm<sup>3</sup>). Finally, a RAVENS map was calculated for each subject. The RAVENS methodology precisely quantifies the volume of tissue in each region of the brain. The process is based on a volume-preserving spatial transformation that ensures that no volumetric information is lost during the process of spatial normalization (Davatzikos et al., 2001).

## 4.1 Selecting the partition model

Following the pre-processing steps, we downsampled the images, cropped them, and obtained images of size  $96 \times 96 \times 96$  mm<sup>3</sup>. We then considered: 1) 64 partitions of size  $24 \times 24 \times 24$  mm<sup>3</sup>; 2) 512 partitions of size  $12 \times 12 \times 12$  mm<sup>3</sup>; and 3) 4096 partitions of size  $6 \times 6 \times 6$  mm<sup>3</sup>. For different values of  $R$ , we selected the number of partitions based on the prediction accuracy of a 10-fold cross validation with the following steps. First, we extracted the features by tensor decomposition for different values of rank  $R$ . Second, to reduce the dimension of the extracted feature matrix, we applied a factor model with fixed rank  $K = 100$ . Third, we run 5000 iterations of the Bayesian probit model with the mixture prior described in Section 2.3 with a burn-in of 3000 samples. Finally we computed the mean prediction accuracy in each model. Results are shown in Table 3.



Table 3: Mean prediction accuracy for a 10-fold cross-validation procedure. There is a smaller error measurement with an increase of the rank  $R$ .

Partition size	$24 \times 24 \times 24$	$12 \times 12 \times 12$	$6 \times 6 \times 6$
$R = 5$	—	—	<b>0.7813</b>
$R = 10$	0.6714	0.7311	0.7613
$R = 20$	0.6937	0.7587	0.7588
$R = 30$	0.5770	0.6544	—
No partition	$1 \times 1 \times 1$		
$R = 100$	0.7498	—	—

## 4.2 Final analysis based on the selected model

Based on the prediction accuracy, we selected the model with partitions of size  $6 \times 6 \times 6$  mm<sup>3</sup> and  $R = 5$ . For the selected model, we run the model described on Section 2.2 with hyperparameters chosen to reflect non-informative priors with  $\nu_{0\tau} = 1$ ,  $\nu_{1\tau} = 10^{-4}$ ,  $\sigma^2 = 10^4$ ,  $\kappa = 10^{-4}$ ,  $\beta_{0\tau_\psi} = 10^{-6}$  and  $\beta_{1\tau_\psi} = 10^{-6}$ . In the first screening procedure, we eliminate the partitions whose features, extracted from the tensor decomposition, were zero since they may be not relevant to predict AD outcome. From the 4096 original partitions, only 1695 passed the first screening. Figure 5 shows the correlation between the features extracted in the first screening step; panel a shows the correlations for all 8475 features and panel b gives a zoom in the first 200 correlations. Inspecting the figure, we observe a very high correlation of features within the partitions and between nearby partitions. That justifies the need of the factor model in Equation (5).

Finally, we ran the Gibbs sampler algorithm described in Section 2.4 for 5,000 iterations with a burn-in period of 3,000 iterations. First, we compute a 95% credible interval for the coefficients  $\mathbf{b}$  and conclude 36 features are important to predict AD outcome. Figure 6 shows the 5 most important bases projected into the image space. The importance is given by the absolute value of the posterior mean in each one of the 36 selected features.

Second, we let  $\tilde{P} = \mathbf{D}\mathbf{b}$ , we computed the posterior mean of the projection  $\mathcal{P} = \|\Lambda; \mathbf{A}^{(1)}, \mathbf{A}^{(2)}, \mathbf{A}^{(3)}, \tilde{P}\|$ . In addition, we compute a 95% credible interval for  $\mathcal{P}$ . Figure 7 shows the results. Inspecting Figure 7 reveals the regions of the biomarkers selected to predict the AD outcome.

To find out specific locations in the brain responsible for predicting AD outcome, we label the locations indicated on the right panel of Figure 7 considering the Jülich atlas (Eickhoff et al., 2005). The largest biomarker is found in the white matter region known as cingulum, specifically in the posterior cingulate cortex, as shown in Table 4, Appendix

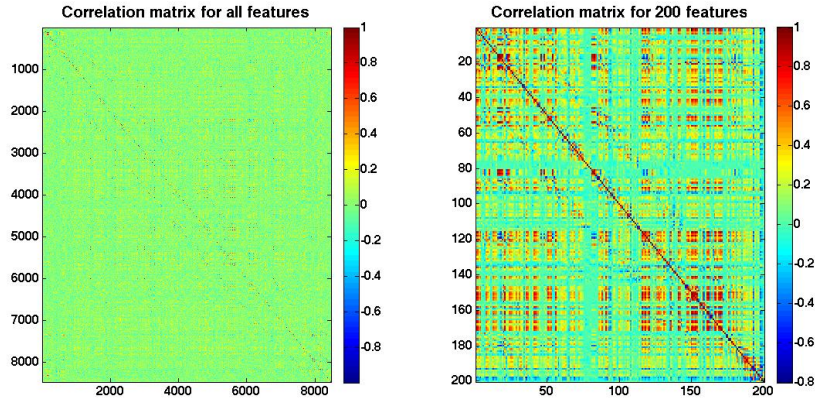


Figure 5: Left panel shows the results for the correlation of the columns of the matrix of features  $L$  obtained in a first screening procedure. The right panel shows the same graphic with the first 200 features. We observe a high correlation between features of the same partitions and neighboring partitions.




Figure 6: Panel (a) shows a 3D rendering of the Ravens map of a random selected control patient. Panels (b)-(f) show a 3D rendering of the 5 most important bases projected into the image space. The importance is given by the absolute value of the posterior mean in each one of the 36 selected features.

A. The posterior cingulate cortex is a limbic lobe that seem to be involved early and consistently in AD (Vemuri and Jack Jr, 2010). The next biomarker is a region known by the fornix, which is white-matter tract linking the hippocampus to several subcortical and cortical regions. Given that these areas are important for successful learning and memory, their ability to communicate with one another via the fornix may also be critical for performance in tasks challenging these cognitive domains (Postans et al., 2014). Important biomarkers were also found within the gray matter tissue, including the posterior parietal lobe (Table 4). The parietal lobe has been linked to neuropsychological deficits in AD likely due to its strong connectivity to other brain areas, and to the wide range of cognitive functions relying on parietal lobe functioning (Jacobs et al., 2012). The most prominent biomarkers were the superior parietal cortex (Table 4). This region is important for spatial processing, selective attention, and spatial and non-spatial working memory (Jacobs et al., 2012). Another important biomarker found is the hippocampus which is associated with learning and consolidation of explicit memories from short-term memory to cortical memory storage for the long term (Campbell and MacQueen, 2004). Previous studies have shown that this region is particularly vulnerable to Alzheimer’s disease pathology and already considerably damaged at the time clinical symptoms first appear (Schuff et al., 2009; Braak and Braak, 1998). Other important biomarkers found by TPRM are shown in Table 4, Appendix A.




Figure 7: Left panel shows the results for the absolute value of the projection  $\mathcal{P}$  for the ADNI dataset. Colors on the right side of the colorbar indicate regions where differences are higher between the control group and the group diagnosed with Alzheimer’s disease. The panel on the right shows a 95% credible interval for  $\mathcal{P}$ . The colored parts indicate the regions of the biomarkers selected to predict the onset of AD.

## 5 Discussion

We have proposed a Bayesian tensor partition regression model (TPRM) to establish an association between imaging tensor predictors and clinical outcomes. The ultra-high dimensionality of imaging data is dramatically reduced by using the proposed partition model. Our TPRM efficiently addresses some key features of neuroimaging data: relatively low signal to noise ratio, spatially clustered effect regions, and the tensor structure of imaging data.

Our simulation studies showed a great performance of the TPRM when predicting a binary outcome from images in 2D and 3D. The TPRM provided much higher prediction accuracy when compared to both, the traditional functional PCA and a tensor model that does not include partitions. In addition, the TPRM is able to capture regions of interest that differ between two groups in scenarios where the fPCA cannot perform well.

Our application to the ADNI dataset showed that our TPRM is able to efficiently reduce and identify relevant biomarkers to predict Alzheimer’s disease outcome, overcoming challenges posed by the complexity of the imaging data such as ultra-high dimensionality and multicollinearity of features.

## A Data analysis supporting material

Table 4: Biomarkers that are relevant to predict AD outcome, based on the Jülich atlas.

Region	# of significant voxels
WM Cingulum R	6547
WM Fornix	4193
GM Premotor cortex BA6 R	4186
WM Superior longitudinal fascicle R	3276
GM Visual cortex V1 BA17 R	2449
GM Superior parietal lobule 7M L	1896
GM Visual cortex V2 BA18 L	1549
GM Superior parietal lobule 7A R	1454
GM Visual cortex V2 BA18 R	1440
GM Primary somatosensory cortex BA3a R	1436
GM Inferior parietal lobule PFm R	1310
WM Corticospinal tract L	1299
GM Inferior parietal lobule PF R	1062
GM Superior parietal lobule 7P R	922
GM Primary motor cortex BA4p R	908
GM Superior parietal lobule 5Ci R	900
WM Superior occipito-frontal fascicle R	889
GM Hippocampus cornu ammonis R	827
WM Acoustic radiation R	807

## B Gibbs sampling algorithm for TPRM

We provide the Gibbs sampling algorithm to sample from the posterior distribution (10) in Section 2.4. It involves sampling from a series of conditional distributions, while each of the modeling components is updated in turn. As an illustration, we divide the whole image into  $S$  equal sized regions and assume  $y_i \sim \text{Bernoulli}(\mu_i)$  with the link function  $h(\cdot)$  being the probit function. By following Albert and Chib (1993), we introduce a normally distributed latent variable,  $w_i$ , such that  $w_i \sim N(\mu_i, 1)$ ;  $y_i = \mathbf{1}(w_i > 0)$ , where  $\mathbf{1}(\cdot)$  is an indicator function of an event.

The complete Gibbs sampler algorithm proceeds as follows.

(a.0) Generate  $\mathbf{w} = (w_1, \dots, w_n)^T$  from

$$w_i | y_i = 0 \sim \mathbf{1}(w_i \leq 0) \text{N}(\mathbf{z}_i^T \boldsymbol{\gamma} + \sum_{s=1}^S \mathbf{g}_i^{(s)T} \mathbf{b}^{(s)}, 1),$$

$$w_i | y_i = 1 \sim \mathbf{1}(w_i \geq 0) \text{N}(\mathbf{z}_i^T \boldsymbol{\gamma} + \sum_{s=1}^S \mathbf{g}_i^{(s)T} \mathbf{b}^{(s)}, 1).$$

(a.1) Update  $\tau(s)$  from its full conditional distribution

$$\tau(s) | \dots \sim \text{Gamma}(\nu_{0\tau} + (N \prod_{d=1}^D p_d)/2, \nu_{1\tau} + (1/2) \sum_{i,j_1,\dots,j_D} (x_{ij_1,\dots,j_D}^*(s))^2),$$

where  $x_{ij_1,\dots,j_D}^*(s) = \{\mathcal{X}^{(s)} - \|\Lambda^{(s)}; \mathbf{A}_s^{(1)}, \mathbf{A}_s^{(2)}, \dots, \mathbf{A}_s^{(D)}, \mathbf{L}^{(s)}\|\}_{ij_1,\dots,j_D}$

(a.2) Update  $\{\mathbf{A}_s^{(d)}\}_{j_d r}$  from its full conditional distribution given by

$$\{\mathbf{A}_s^{(d)}\}_{j_d r} | \dots \sim \text{N} \left( \frac{\tau(s) \langle \widehat{\mathcal{X}}_{(-r)}^{s(j_d)}, \mathcal{I}_{(-d)}^s \rangle}{\tau(s) \langle \mathcal{I}_{(-d)}^s, \mathcal{I}_{(-d)}^s \rangle + p_d}, (\tau(s) \langle \mathcal{I}_{(-d)}^s, \mathcal{I}_{(-d)}^s \rangle + p_d)^{-1} \right),$$

where  $\mathcal{I}_{(-d)}^s = \|\Lambda^{(s)}; \mathbf{A}_s^{(1)}, \dots, \mathbf{A}_s^{(d-1)}, \mathbf{A}_s^{(d+1)}, \dots, \mathbf{A}_s^{(D)}, \mathbf{L}^{(s)}\|$ ,  $\widehat{\mathcal{X}}_{(-r)}^s$  is given by

$$\mathcal{X}^{(s)} - \|\Lambda^{(s)}; \mathbf{A}_s^{(1)}, \mathbf{A}_s^{(2)}, \dots, \mathbf{A}_s^{(D)}, \mathbf{L}_i^{(s)}\| + \|\Lambda^{(s)}; \{\mathbf{A}_s^{(1)}\}_{:,r}, \{\mathbf{A}_s^{(2)}\}_{:,r}, \dots, \{\mathbf{A}_s^{(D)}\}_{:,r}, \{\mathbf{L}_i^{(s)}\}_{:,r}\|,$$

and  $\widehat{\mathcal{X}}_{(-r)}^{s(j_d)}$  is a subtensor fixed at the entry  $j_d$  along the  $d$ -th dimension of  $\widehat{\mathcal{X}}_{(-r)}^s$ .

(a.3) Update  $\{\mathbf{L}_s\}_{ir}$  from its full conditional distribution given by

$$\{\mathbf{L}_s\}_{ir} | \dots \sim \text{N} \left( \frac{\tau(s) \langle \widehat{\mathcal{X}}_{(-r)}^{s(i)}, \mathcal{I}^s \rangle}{\tau(s) \langle \mathcal{I}^s, \mathcal{I}^s \rangle + N}, (\tau(s) \langle \mathcal{I}^s, \mathcal{I}^s \rangle + N)^{-1} \right),$$

where  $\mathcal{I}^s = \|\Lambda^{(s)}; \mathbf{A}_s^{(1)}, \dots, \mathbf{A}_s^{(D)}\|$ ,  $\widehat{\mathcal{X}}_{(-r)}^s$  is the same as above, and  $\widehat{\mathcal{X}}_{(-r)}^{s(i)}$  is a subtensor fixed at the  $i$ -th entry along the subject dimension of  $\widehat{\mathcal{X}}_{(-r)}^s$ .

(a.4) Update  $\Lambda^{(s)}$  from its full conditional distribution

$$\lambda_r^{(s)} | \dots \sim \text{N} \left( \frac{\tau(s) \langle \widehat{\mathcal{X}}_{(-r)}^s, \mathcal{L}^s \rangle}{\tau(s) \langle \mathcal{L}^s, \mathcal{L}^s \rangle + \kappa}, (\tau(s) \langle \mathcal{L}^s, \mathcal{L}^s \rangle + \kappa)^{-1} \right),$$

where  $\mathcal{L}^s = \|\mathbf{1}_R; \mathbf{A}_s^{(1)}, \dots, \mathbf{A}_s^{(D)}, \mathbf{L}^{(s)}\|$  and  $\mathbf{1}_R$  is a vector of ones of size  $R$ .

(a.5) Update  $\mathbf{g}_k$  from its full conditional distribution

$$\mathbf{g}_k | - \sim \text{N}(\mu_g, \Sigma_g), \quad \Sigma_g = \left( n\mathbf{I}_n + \tau_\psi \sum_{j=1}^{P_L} d_{kj}^2 \right)^{-1} \quad \text{and} \quad \mu_g = \tau_\psi \Sigma_g \sum_{j=1}^{P_L} d_{kj} \mathbf{l}_j^{*-k},$$

where  $\mathbf{l}_j^{*-k} = \mathbf{L} - \mathbf{G}\mathbf{d}_j + d_{kj}\mathbf{g}_k$  for  $j = 1, \dots, P_L$ .

(a.6) Update  $d_{kj}$  for  $j = 1, \dots, P_L$  from its full conditional distribution

$$d_{kj} | - \sim \text{N} \left( \tau_\psi \Sigma_d \sum_{j=1}^{P_L} \mathbf{g}_k^T \mathbf{l}_j^{*-k}, \Sigma_d \right)$$

where  $\Sigma_d = \left( 1 + \tau_\psi \sum_{j=1}^{P_L} \mathbf{g}_k^T \mathbf{g}_k \right)^{-1}$ .

(a.7) Update  $\psi_{ij}$  from its full conditional distribution

$$\psi_{ij} | - \sim \text{Gamma} \left( \beta_{0\psi} + NP_L/2, \beta_{1\psi} + (\mathbf{L}^{*T} \mathbf{L}^*)/2 \right),$$

where  $\mathbf{L}^* = \mathbf{L} - \mathbf{G}\mathbf{D}$ .

(a.8) Update  $\delta_k$  from its full conditional distribution

$$\delta_k \sim \text{bernoulli}(\tilde{p}_1/\tilde{p}_1 + \tilde{p}_0),$$

where  $\tilde{p}_1 = \pi \exp\{-(1/2\sigma^2)(\mathbf{b}_k)^2\}$  and  $\tilde{p}_0 = \pi \exp\{-(1/2\epsilon)\mathbf{b}_k^2\}$ .

(a.9) Update  $\mathbf{b}$  from its full conditional distribution

$$b_k | \delta_k = 1 \sim \text{N}(\sum_i \tilde{w}_i g_{ik} / \sum_i g_{ik}^2 + 1/\sigma^2, (\sum_i g_{ik}^2 + 1/\sigma^2)^{-1}),$$

$$b_k | \delta_k = 0 \sim \text{N}(\sum_i \tilde{w}_i g_{ik} / \sum_i g_{ik}^2 + 1/\epsilon, (\sum_i g_{ik}^2 + 1/\epsilon)^{-1}),$$

where  $\tilde{w}_i^{(s)} = w_i - \mathbf{z}_i^T \boldsymbol{\gamma} - \sum_{s'=1}^S \mathbf{g}_i^{(s')T} \mathbf{b}^{(s')} + g_{ir}^{(s)T} b_r^{(s)}$ .

(a.10) Update  $\pi$  from its full conditional distribution

$$\pi | \dots \sim \text{beta}(\alpha_{0\pi} + \sum_{s,r} \delta_r^{(s)}, \alpha_{1\pi} + SR - \sum_{s,r} \delta_r^{(s)}).$$

(a.11) Update  $\gamma$  from its full conditional distribution

$$\gamma | \dots \sim N \left( \Sigma_{\gamma}^{*-1} \left( v\gamma^* + \mathbf{Z}^T \mathbf{w}_{\gamma}^* \right), \Sigma_{\gamma}^{*-1} \right),$$

where  $\Sigma_{\gamma}^* = v\mathbf{I}_q + \mathbf{Z}^T \mathbf{Z}$  and  $\mathbf{w}_{\gamma}^* = \mathbf{w} - \mathbf{G}^T \mathbf{b}$ .

(a.12) Update  $v$  from its full conditional distribution

$$v | \dots \sim \text{Gamma} \left( \nu_{0v} + q/2, \nu_{1v} + (\gamma^T \gamma)/2 \right).$$

All the tensor operations described in steps (a.1) – (a.4) can be easily computed using Bader et al. (2012), <http://www.sandia.gov/~tgkolda/TensorToolbox/index-2.5.html>.

## Acknowledgements

Data collection and sharing for this project was funded by the Alzheimer’s Disease Neuroimaging Initiative (ADNI) (National Institutes of Health Grant U01 AG024904) and DOD ADNI (Department of Defense award number W81XWH-12-2-0012). ADNI is funded by the National Institute on Aging, the National Institute of Biomedical Imaging and Bioengineering, and through generous contributions from the following: Alzheimer’s Association; Alzheimer’s Drug Discovery Foundation; Araclon Biotech; BioClinica, Inc.; Biogen Idec Inc.; Bristol-Myers Squibb Company; Eisai Inc.; Elan Pharmaceuticals, Inc.; Eli Lilly and Company; EuroImmun; F. Hoffmann-La Roche Ltd and its affiliated company Genentech, Inc.; Fujirebio; GE Healthcare; ; IXICO Ltd.; Janssen Alzheimer Immunotherapy Research & Development, LLC.; Johnson & Johnson Pharmaceutical Research & Development LLC.; Medpace, Inc.; Merck & Co., Inc.; Meso Scale Diagnostics, LLC.; NeuroRx Research; Neurotrack Technologies; Novartis Pharmaceuticals Corporation; Pfizer Inc.; Piramal Imaging; Servier; Synarc Inc.; and Takeda Pharmaceutical Company. The Canadian Institutes of Health Research is providing funds to support ADNI clinical sites in Canada. Private sector contributions are facilitated by the Foundation for the National Institutes of Health ([www.fnih.org](http://www.fnih.org)). The grantee organization is the Northern California Institute for Research and Education, and the study is coordinated by the Alzheimer’s Disease Cooperative Study at the University of California, San Diego. ADNI data are disseminated by the Laboratory for Neuro Imaging at the University of Southern California.



## References

- Albert, J. H. and Chib, S. (1993). Bayesian analysis of binary and polychotomous response data. *Journal of the American Statistical Association*, 88(422):pp. 669–679.
- Bader, B. W., Kolda, T. G., et al. (2012). Matlab tensor toolbox version 2.5. Available online.
- Bair, E., Hastie, T., Paul, D., and Tibshirani, R. (2006). Prediction by supervised principal components. *Journal of the American Statistical Association*, 101:119–137.
- Beckmann, C. F. and Smith, S. M. (2005). Tensorial extensions of independent component analysis for multisubject fMRI analysis. *NeuroImage*, 25(1):294 – 311.
- Bickel, P. and Levina, E. (2004). Some theory for Fisher’s linear discriminant function, ‘naive Bayes’, and some alternatives when there are many more variables than observations. *Bernoulli*, 10:989–1010.
- Braak, H. and Braak, E. (1998). Evolution of neuronal changes in the course of Alzheimer’s disease. In Jellinger, K., Fazekas, F., and Windisch, M., editors, *Ageing and Dementia*, volume 53 of *Journal of Neural Transmission. Supplementa*, pages 127–140. Springer Vienna.
- Breiman, L., Friedman, J., Olshen, R., and Stone, C. (1984). *Classification and Regression Trees*. Wadsworth, California.
- Caffo, B., Crainiceanu, C., Verduzco, G., Joel, S., S.H., M., Bassett, S., and Pekar, J. (2010). Two-stage decompositions for the analysis of functional connectivity for fMRI with application to Alzheimer’s disease risk. *Neuroimage*, 51(3):1140–1149.
- Campbell, S. and MacQueen, G. (2004). The role of the hippocampus in the pathophysiology of major depression. *Journal of Psychiatry and Neuroscience*, 29(6):417—426.
- Davatzikos, C., Genc, A., Xu, D., and Resnick, S. M. (2001). Voxel-based morphometry using the RAVENS maps: Methods and validation using simulated longitudinal atrophy. *NeuroImage*, 14(6):1361 – 1369.
- Eickhoff, S. B., Stephan, K. E., Mohlberg, H., Grefkes, C., Fink, G. R., Amunts, K., and Zilles, K. (2005). A new SPM toolbox for combining probabilistic cytoarchitectonic maps and functional imaging data. *NeuroImage*, 25(4):1325 – 1335.

- Fan, J. and Fan, Y. (2008). High-dimensional classification using features annealed independence rules. *Annals of Statistics*, 36:2605–2637.
- Friedman, J. (1991). Multivariate adaptive regression splines (with discussion). *Annals of Statistics*, 19:1–141.
- George, E. I. and McCulloch, R. E. (1993). Variable selection via Gibbs sampling. *Journal of the American Statistical Association*, 88(423):pp. 881–889.
- George, E. I. and McCulloch, R. E. (1997). Approaches for Bayesian variable selection. *Statistica Sinica*, 7:339–373.
- Gonçalves, F., Gamerman, D., and Soares, T. (2013). Simultaneous multifactor DIF analysis and detection in item response theory. *Computational Statistics & Data Analysis*, 59(0):144 – 160.
- Hastie, T., Tibshirani, R., and Friedman, J. (2009). *The Elements of Statistical Learning: Data Mining, Inference, and Prediction (2nd)*. Springer, Hoboken, New Jersey.
- Huang, M., Nichols, T., Huang, C., Yang, Y., Lu, Z., Feng, Q., Knickmeyer, R. C., Zhu, H., , and for the Alzheimer’s Disease Neuroimaging Initiative (2014). FVGWAS: Fast voxelwise genome wide association analysis of large-scale imaging genetic data. *Preprint*.
- Jacobs, H. I., Boxtel, M. P. V., Jolles, J., Verhey, F. R., and Uylings, H. B. (2012). Parietal cortex matters in Alzheimer’s disease: An overview of structural, functional and metabolic findings. *Neuroscience & Biobehavioral Reviews*, 36(1):297 – 309.
- Johnstone, I. M. and Lu, A. Y. (2009). On consistency and sparsity for principal components analysis in high dimensions. *Journal of the American Statistical Association*, 104:682–693.
- Kolda, T. G. and Bader, B. W. (2009). Tensor decompositions and applications. *SIAM Rev.*, 51(3):455–500.
- Krishnan, A., Williams, L., McIntosh, A., and Abdi, H. (2011). Partial least squares (PLS) methods for neuroimaging: a tutorial and review. *Neuroimage*, 56:455–475.
- Martinez, E., Valdes, P., Miwakeichi, F., Goldman, R. I., and Cohen, M. S. (2004). Concurrent EEG/fMRI analysis by multiway partial least squares. *NeuroImage*, 22(3):1023 – 1034.

- Mayrink, V. D. and Lucas, J. E. (2013). Sparse latent factor models with interactions: Analysis of gene expression data. *The Annals of Applied Statistics*, 7(2):799–822.
- Mitchell, T. J. and Beauchamp, J. J. (1988). Bayesian variable selection in linear regression. *Journal of the American Statistical Association*, 83(404):pp. 1023–1032.
- Müller, H.-G. and Yao, F. (2008). Functional additive models. *Journal of the American Statistical Association*, 103(484):1534–1544.
- Postans, M., Hodgetts, C. J., Mundy, M. E., Jones, D. K., Lawrence, A. D., and Graham, K. S. (2014). Interindividual variation in fornix microstructure and macrostructure is related to visual discrimination accuracy for scenes but not faces. *The Journal of Neuroscience*, 34(36):12121–12126.
- Ramsay, J. O. and Silverman, B. W. (2005). *Functional data analysis*. Springer Series in Statistics. Springer, New York, second edition.
- Reiss, P. T. and Ogden, R. T. (2010). Functional generalized linear models with images as predictors. *Biometrics*, 66(1):61–69.
- Schuff, N., Woerner, N., Boreta, L., Kornfield, T., Shaw, L. M., Trojanowski, J. Q., Thompson, P. M., Jack Jr, C. R., and Weiner, M. W. (2009). MRI of hippocampal volume loss in early Alzheimer’s disease in relation to ApoE genotype and biomarkers. *Brain*, 132(4):1067–1077.
- Tibshirani, R., Hastie, T., Narasimhan, B., and Chu, G. (2002). Diagnosis of multiple cancer types by shrunken centroids of gene expression. *Proceedings of the National Academy of Sciences*, 99:6567–6572.
- Vemuri, P. and Jack Jr, C. R. (2010). Role of structural MRI in Alzheimer’s disease. *Alzheimer’s Research and Therapy*, pages 2–23.
- Xinghao Ding, L. and Carin, L. (2011). Bayesian robust principal component analysis. *Imaging Processing, IEEE Transactions on*, 20(12):3419–3430.
- Zhang, H. P. and Singer, B. H. (2010). *Recursive Partitioning and Applications (2nd)*. Springer, New York.
- Zhou, H., Li, L., and Zhu, H. (2013). Tensor regression with applications in neuroimaging data analysis. *Journal of the American Statistical Association*, 108(502):540–552.